

KCGIS Software Migration Post Mortem

January 2007

Project Overview

The overall goal of the King County GIS software migration project was to migrate county GIS software from ArcInfo 7.1.2 / ArcView 3.x to ArcGIS 8.3. Over the course of the migration, ESRI has released three subsequent versions of the software. As of the end of the Software Migration project, the current ArcGIS software version is 9.2 (released November 2006, with a major patch issued in January 2007).

The magnitude of this software migration is analogous to moving from driving a 1963 VW bus on a dirt road to driving a 2000 Accord on I-5. Not only is the equipment different, but the whole concept of driving has changed. Likewise, King County GIS has had to not only learn, but embrace an entirely new way of doing business with GIS: an entirely new way of collecting, editing, storing and maintaining data; new and exciting tools within the core GIS software; and a radical new approach of building our own GIS applications for our data maintainers and users. The new concepts at the center of this migration will allow us to use GIS in ways we could not have imagined as little as three years ago, when we started the project.

Within that framework, our specific goals were straightforward: migrate our enterprise and agency data from coverage to geodatabase maintenance; give users the necessary tools to be successful in the new environment via training and enterprise applications; minimize licensing costs; and identify and mitigate agency concerns, while causing as little disruption as possible to the day-to-day and overall business of agencies.

Analysis of Objectives and Milestones:

There were four main objectives of the GIS software migration. With all of these objectives met, the migration is considered complete:

- The primary data warehouse (except cadastral base) is the read-only SDE geodatabase. The shapefile library is maintained for "legacy" ArcView 3.x users. The coverage library no longer exists.

Objective MET.

- Data editing and posting takes place in the SDE geodatabase environment.

Objective MET. Editing takes place in whatever format the user desires, but posting is via SDE. There has been good progress in stewards establishing good workflow editing from persistent PGDB or structured editing environments in SDE databases.

- Enterprise applications are in place to facilitate data access, management, and editing where necessary.

Objective MET.

Agency-specific applications are migrated or their relevant functionality integrated into other business applications.

Remaining agency applications that will break when WILDFIRE goes offline either have workarounds in place until effective long-term replacement can be deployed, or will be allowed to break. Other outstanding issues will be dealt with as needed. **Objective Met.**

- End users have been categorized and trained, and have either migrated to ArcGIS, ArcIMS, or have been declared as "legacy" ArcView 3.x users.

Objective MET. This verbiage was changed to reflect the current ongoing iterative process. User training "In Progress" is acceptable for completion.

To track progress toward these objectives, seven major milestones were identified.

- Develop training curriculum. Sort all GIS users into categories. Develop a training curriculum for each user category, using available, cost-effective, and appropriate courses from ESRI classroom, ESRI Virtual Campus, KCGIS Center courses and modules, and other sources. **Done.**
- Complete preliminary data review. Conduct a fitness review of every internally-maintained coverage in the current GIS data warehouse (/plibrary). Layers that do not pass review should be archived and deleted immediately. **Done.**
- Create agency migration plans. Categorize agency business and technical needs into functional groups and prioritize based on common needs. Use this information and that acquired from agency needs assessment, data design, and geodatabase design and implementation to create a migration plan for each agency. **Determined superfluous, and unneeded extra workload.**
- Implement prototype SDE production geodatabase. Implement and test a prototype enterprise SDE production geodatabase, using copies of core data layers. Ensure that stewards can connect to their data, edit it, and publish edited data to the data warehouse. Devise and publish methodology and appropriate guidelines for stewards,

developers and analysts. Note that this will not include the cadastral data model, but will assume the presence of the parcel layer. **Done.**

- Optimize and migrate internally-maintained data to the production geodatabase. Determine layer dependencies and prioritize layers and layer groups to migrate based on agency needs. Design, implement, and test optimization processes based on a set of prototype layers. Optimize and migrate data. Remove migrated data from /plibrary. **Done.**
- Migrate front-end enterprise applications for data access and management. For each application included: determine need, design, implement, test and deploy. Create and publish user documentation.
 - Data posting, metadata: **Done.**
 - Data Access: **Postponed, pending further review.** The Application Development group determined (with input from users) that the replacement of existing ArcView data access applications is less important as off-the-shelf ArcGIS handles access well enough to deprioritize this task.
 - Easy Map Creation: **In Progress.** An ArcGIS addon will be finalized and deployed in the first part of 2007.
 - Access to RECDNET: **Postponed, pending further review.** ArcView users will still be able to access the Assessor tables via ParcelTools, but RECDNET will be offline in January 2007. The Quarter section map functionality has been replaced by the easily-accessible maps on the Assessor web site. This means that users will need to upgrade to ArcGIS 9.X to view live data and make maps using that data.
- Migrate users. For each user (or group of users, depending on the agency), determine the best migration path then implement. **Mostly Done.** Implementation is up to individual agencies, with help as needed from the Software Migration Workgroup. At this time, users have either been migrated, are planning to migrate, or will not migrate.

Identification of Ongoing Work:

Data:

- Creation of Data Management & Coordination (DMC) group (KCGIS Center Data Coordinator, lead). This group will meet bimonthly for the foreseeable future. It will focus on resolving specific data issues and also address cross-agency database configuration topics.
- All data objects (layers and tables) that were represented in Plibrary as coverages or in Plibrary2 as shapefiles have been migrated to SDE featureclass or SDE tables, save for a few exceptions that are being addressed by the DMC group.
- Work is continuing on creating metadata for all tables. Coming updates to PostRep to support metadata propagation during all posting cycles will enable metadata to be published consistently in all formats.
- A group will soon be evaluating available database capacity against remaining raster datasets to determine how much of the lower priority raster library will be imported into SDE.

Training:

- The Training Workgroup will continue to develop curriculum, guided by and subject to the approval of the GIS Technical Committee.
- This group will continue to meet on a regular basis for the foreseeable future.

Enterprise Application Development:

- The remaining enterprise application tasks (data access and access to Assessments data) have been included in the 2007 KCGIS Priority Work Initiatives.
- The KCGIS Center will continue its lead role in the development and implementation of enterprise applications. The GIS Application Developers group will continue its coordination efforts, including identification of enterprise applications, requirements gathering and general design and implementation plans.
- The AppDev group will continue to meet on a regular basis for the foreseeable future.

Licensing:

- The outstanding licensing tasks (license consolidation) have been taken on by the KCGIS Center; the Licensing Workgroup still exists, but does not have a regular meeting schedule.

Agency Tasks:

- DDES:
 - The counter application is reliant on shapefiles.

What Went Right

We had a high degree of enterprise-wide coordination. This was the single "make or break" factor in the entire project. Staff from KC agencies and the KCGIS Center translated their desire to succeed into an environment of communication, cooperation, and shouldering tasks for the good of the order. Without the efforts, talent, and support of key players across the enterprise, the migration would not have succeeded.

We made the effort to create a comprehensive, well-thought out plan. We first identified the broad areas of interest (data, training, applications, software), as well as the areas of impact (agencies and the enterprise; servers, workstations and the web), and then did our research. Stakeholders were identified and brought into the process at the beginning of the planning stage, so that their needs could be enumerated and addressed. We hewed closely to the migration task list, and changed it when needed.

We stuck to the plan, but we were flexible enough to deal with major changes in thought. The necessary change in our core database design did not throw the project significantly behind. Other, more minor, changes to the plan were dealt with as they came up.

Agencies moved at their own pace, and the early adopters assisted the rest. Agencies have varying business needs and internal resources available for migration. The help provided by the early adopters via documentation and direct assistance was invaluable.

We paid close attention to the details, while keeping the overall goal firmly in mind. This was of critical importance to the data migration. Data stewards looked at every coverage, every table, and every field in each, and constantly asked the question "how does this fit?" And as part of the process, they updated or created metadata. The result is a well-understood, well-documented spatial database with little (but expected) redundancy, and no extraneous layers.

What Went Wrong

Our primary vendor was flaky. Support from ESRI was inconsistent at best. Most of what we learned was through our own efforts and perseverance, in spite of ESRI's negligence in disclosing data issues certain to cause problems for their user community.

ESRI database design did not meet our needs. ESRI's data model and experience is too small in scale and too simplistic in organizational design. The server design assumes centralized database control, which is not effective or desired in organizations as large and complex as KCGIS. Multiple databases quickly run against technical limitations, which we have yet to resolve. Permission options for database management and editing scenarios are oversimplified and inadequate.

At every step, we were forced to come up with our own workarounds, or in some cases, live with the inadequacy inherent to the ESRI model. Ironically, near-future versions of ArcGIS that include true data replication may finally deliver a realistic, workable database solution, nearly two years after we designed and implemented our own.

We had to change our entire data design. After we'd put work into designing a database that would meet nearly all the business needs of the enterprise, we discovered that our data design was technically infeasible, even within the constraints outlined above. We had to regroup, find the next best solution and make it work.

Conclusions

Software migration is an ongoing process. As long as software companies are in business, they will be turning out bigger, more complex (and more expensive) versions of their software and forcing users to upgrade through backward incompatibility, "sunsetting" of support, or inevitable platform retirement. This large migration project was a necessary response to the huge conceptual shift in business process between ArcInfo 7.x and ArcGIS 8/9. And while we do not foresee such a large migration in the future, the fact remains that new versions of ArcGIS are coming out every year, and we'll need to deal with them.

Make a firm plan, but be flexible enough to change its underpinnings without having the whole process come down on your head. When we started the planning process, ESRI was at 8.2, and we intended to migrate to 8.3. By the end of the planning process, ESRI was announcing 9.0. The Software Migration Workgroup discussed whether or not to continue the planned migration to the very stable 8.3, or change gears to migrate to the unstable and buggy 9.0. We decided to keep on track with 8.3. As it turned out, 9.X offers tools necessary for the success of the cadastral conversion (a separate, but highly interdependent project). The first patch released for 9.0 was stable enough that we were able to make the changes necessary in the plan and process to allow us to adapt to the new version.

Slow and steady wins the migration. We probably could have pushed through the migration in less time (the process took nearly three years), but it would have been at the expense of day-to-day business in many (if not all) of the involved agencies. By taking our time, we ensured that all agencies were able to meet their goals, and that necessary data issues were dealt with to everyone's satisfaction. The downside was that stakeholders grew tired of the ongoing process, and interest waned in the last few months.

Pay attention to the details. The big concepts drive the direction and scope of the project, but it is in managing the details that the project will succeed or fail.

Plan well, task appropriately, set reasonable expectations and timelines. The Software Migration group met twice a month for the first 18 or so months; then once a month, then every other month for most the last year as the focus of the group shifted from design and planning to large and small group work sessions and finally to status reporting of offline efforts.